

PostgreSQL Replication Solutions

BRUCE MOMJIAN,
SOFTWARE RESEARCH ASSOCIATES

January, 2004

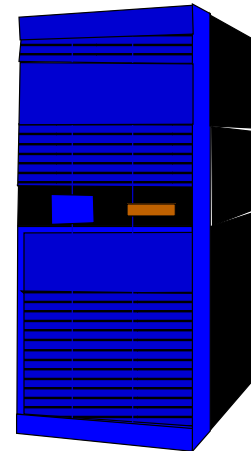
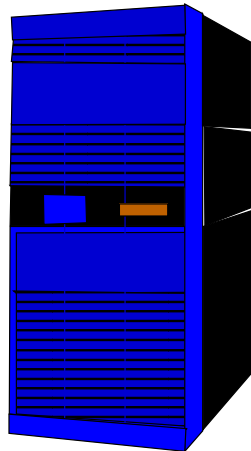
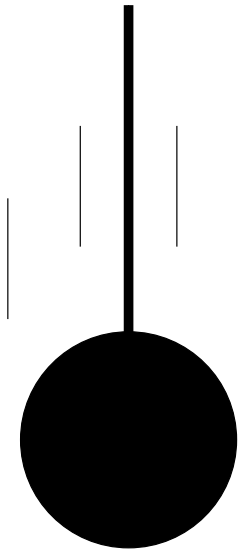


Abstract

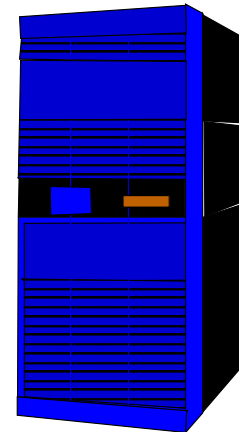
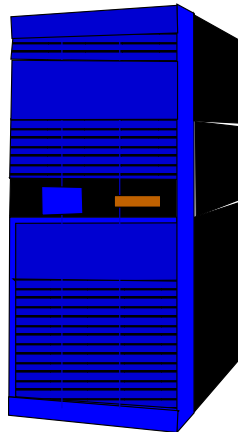
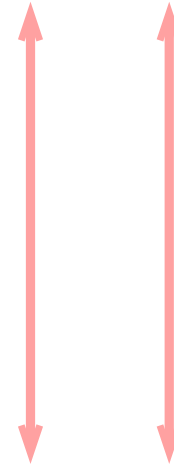
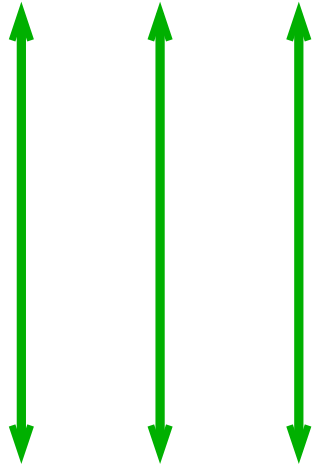
Replication is a complex feature. POSTGRESQL is working toward several replications solutions.

Uses for Replication

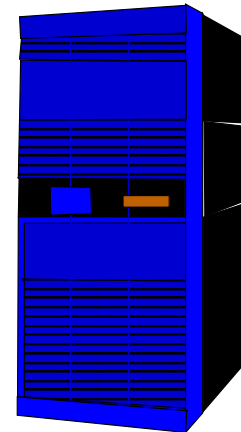
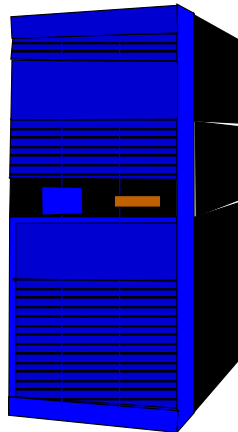
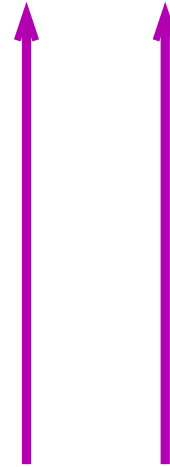
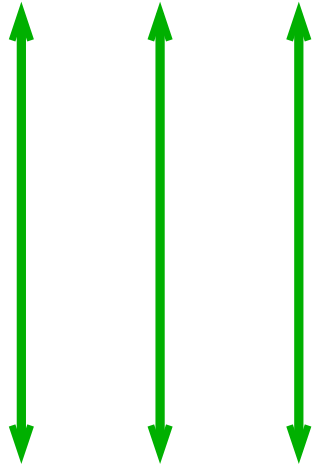
Fail Over



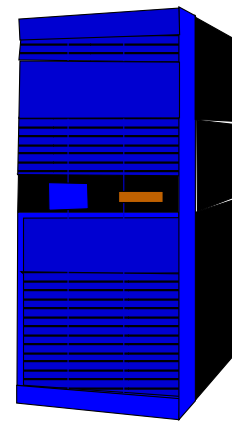
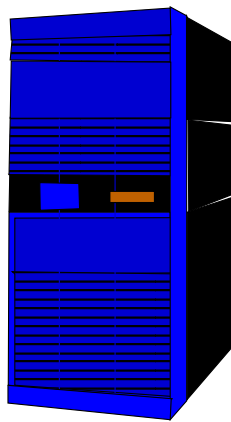
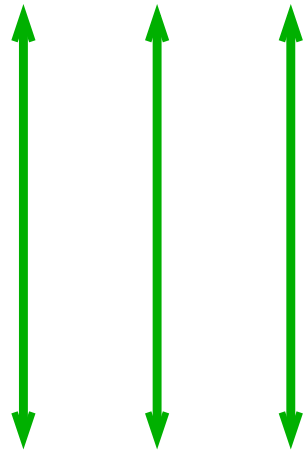
Load Balancing



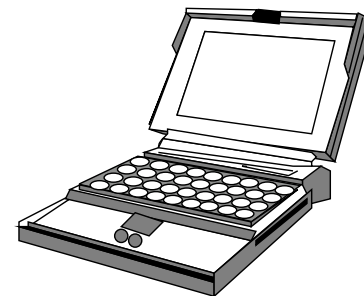
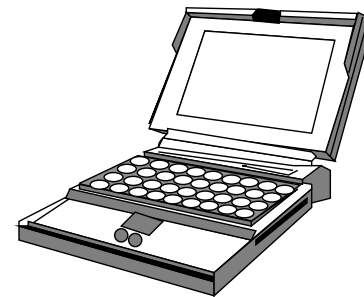
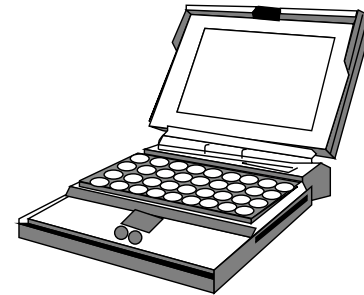
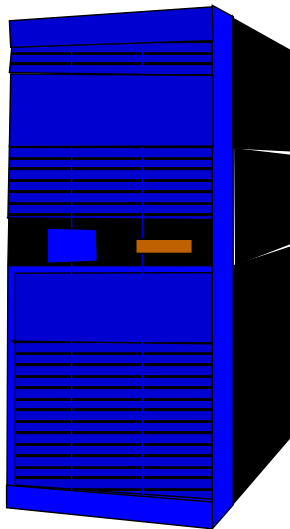
Data Warehousing



Remote Servers

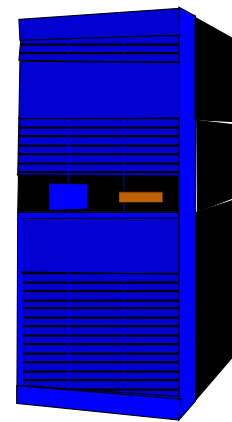
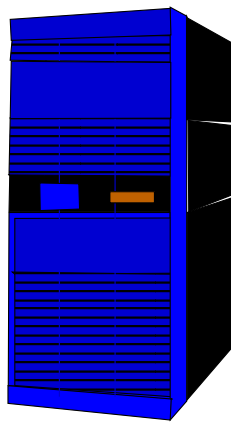
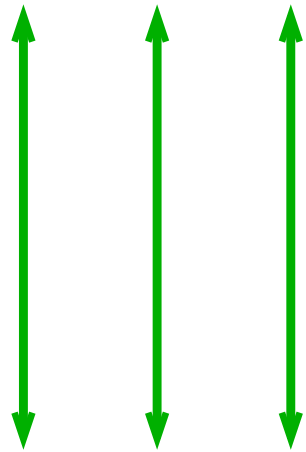


Mobile Servers



Replication Methods

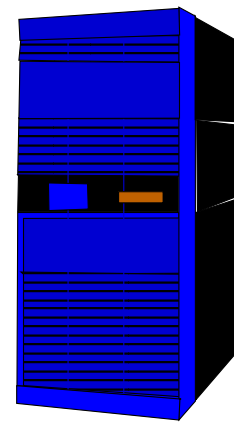
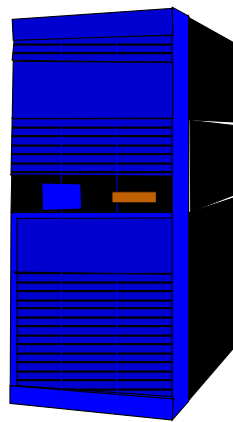
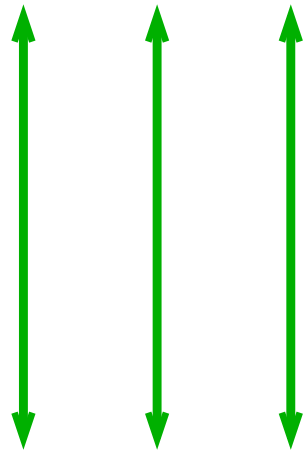
Master/Slave - Asynchronous



Asynchronous



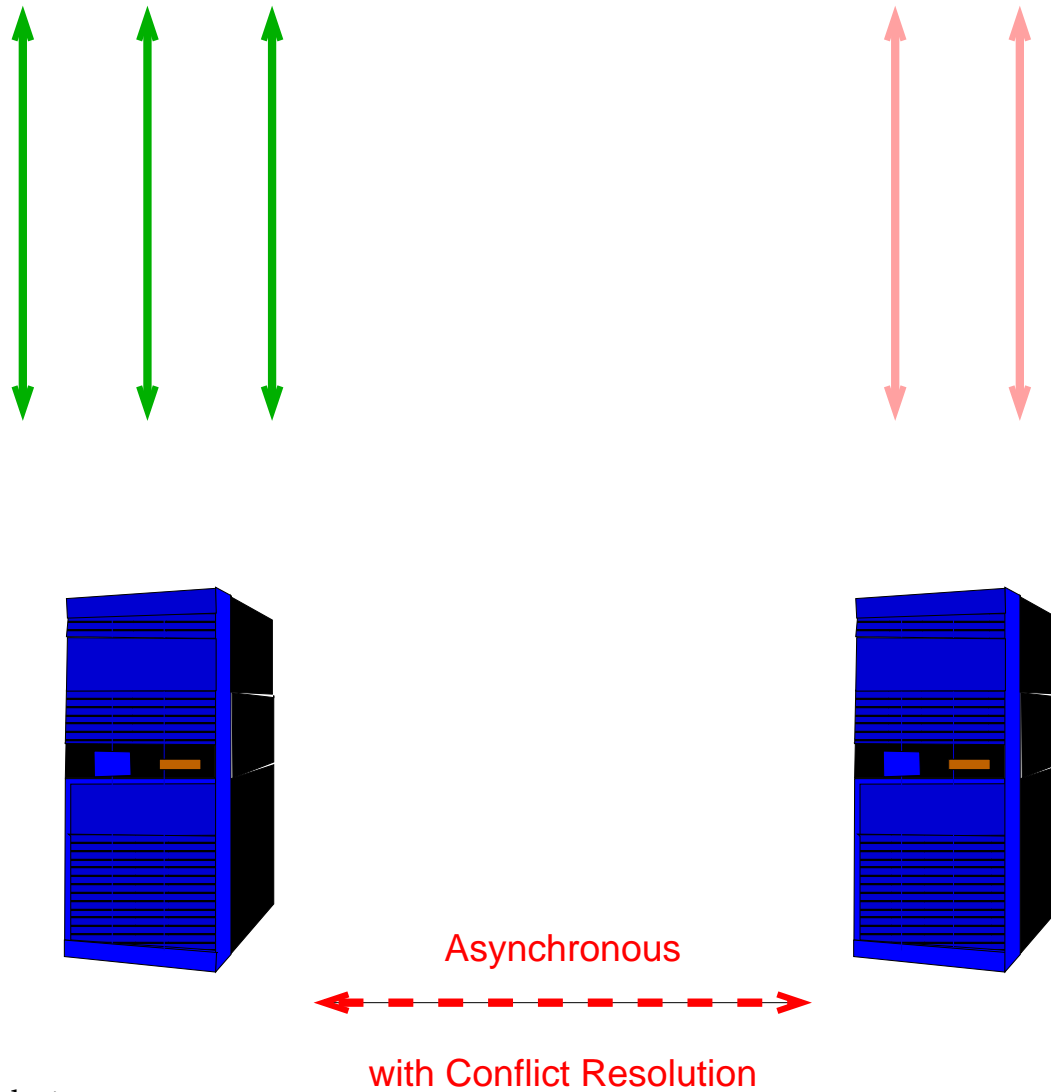
Master/Slave - Synchronous



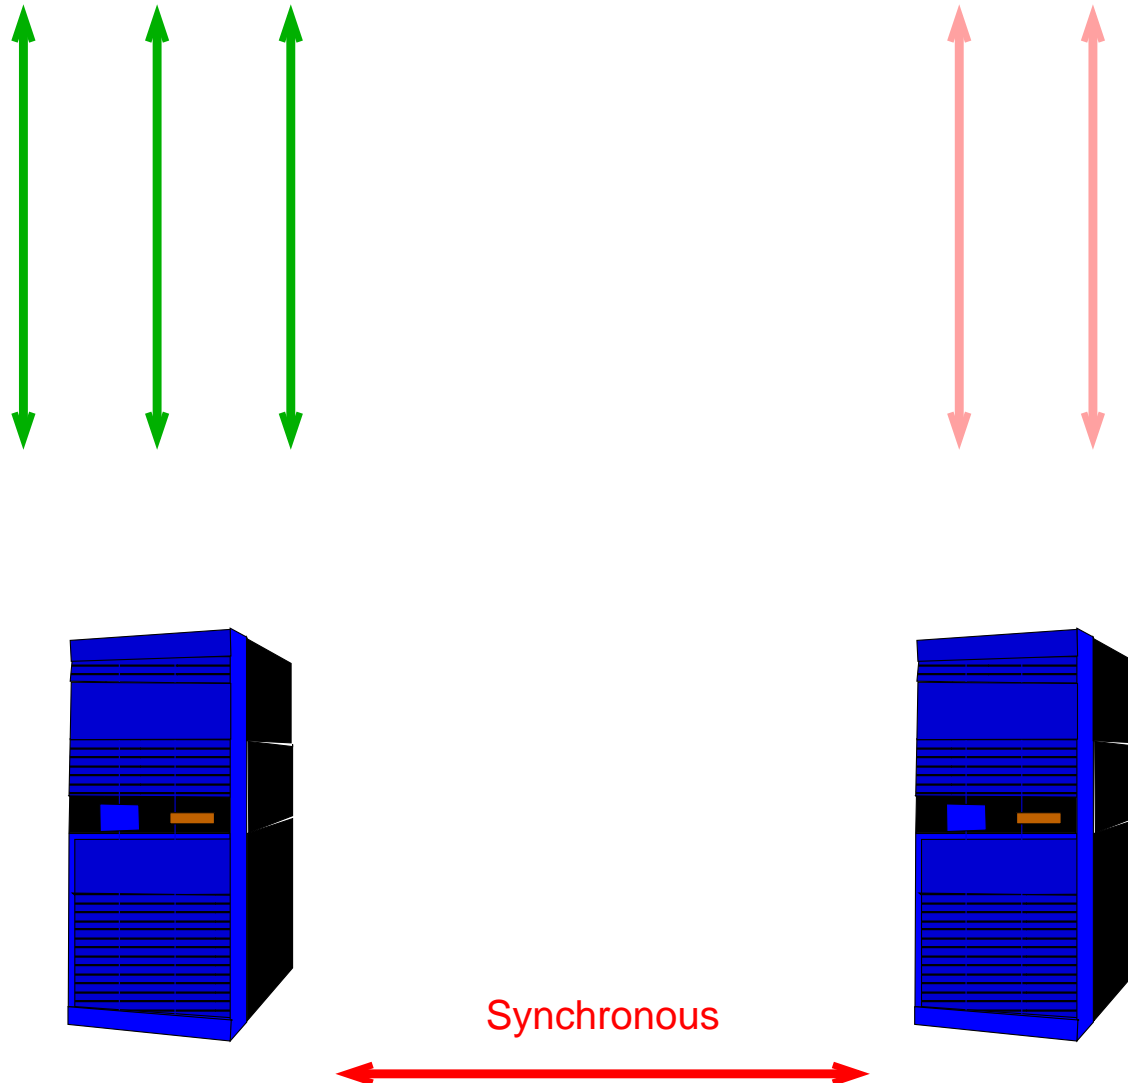
Synchronous



Multi-Master - Asynchronous

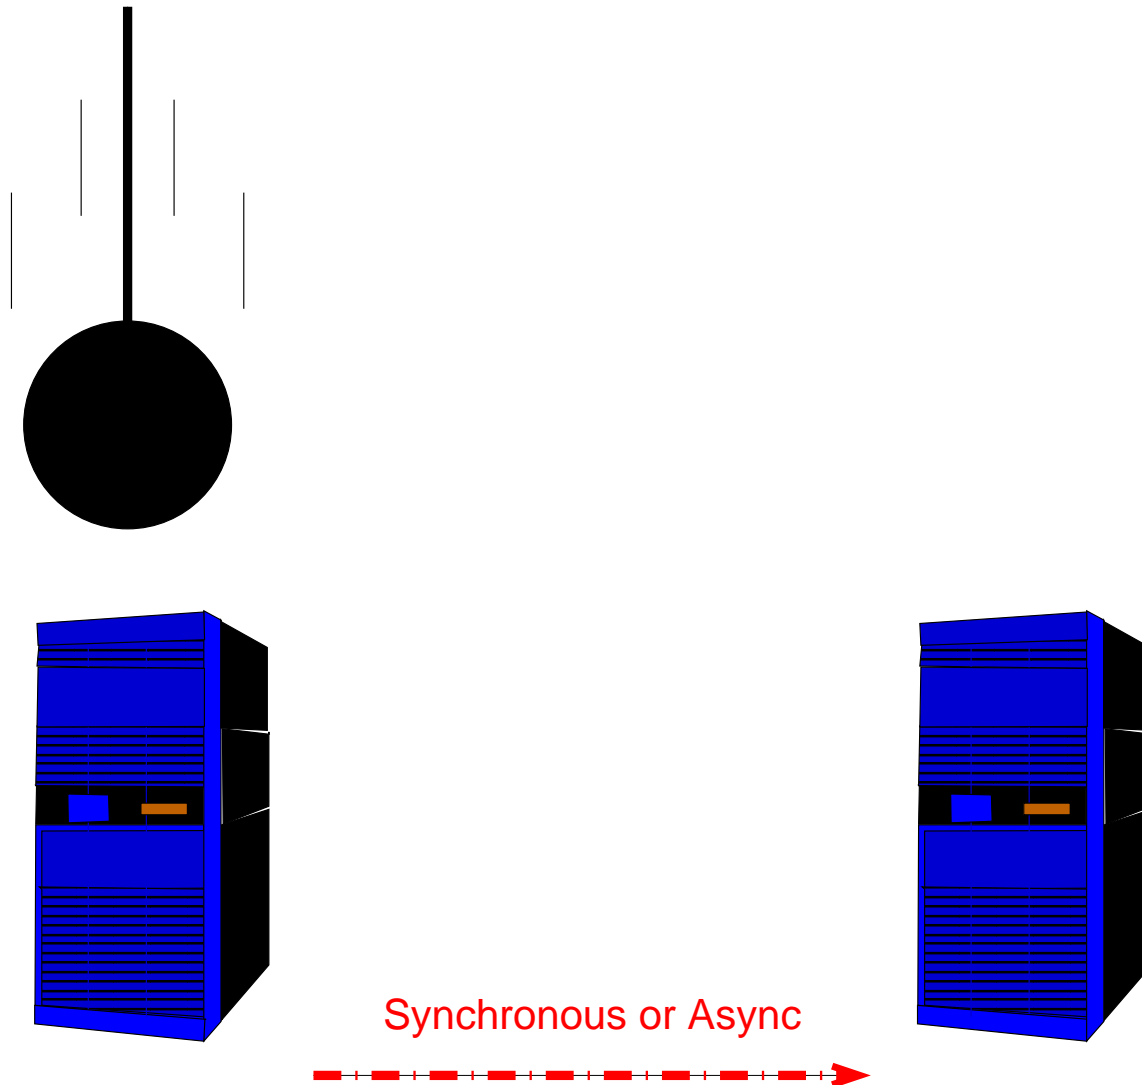


Multi-Master - Synchronous

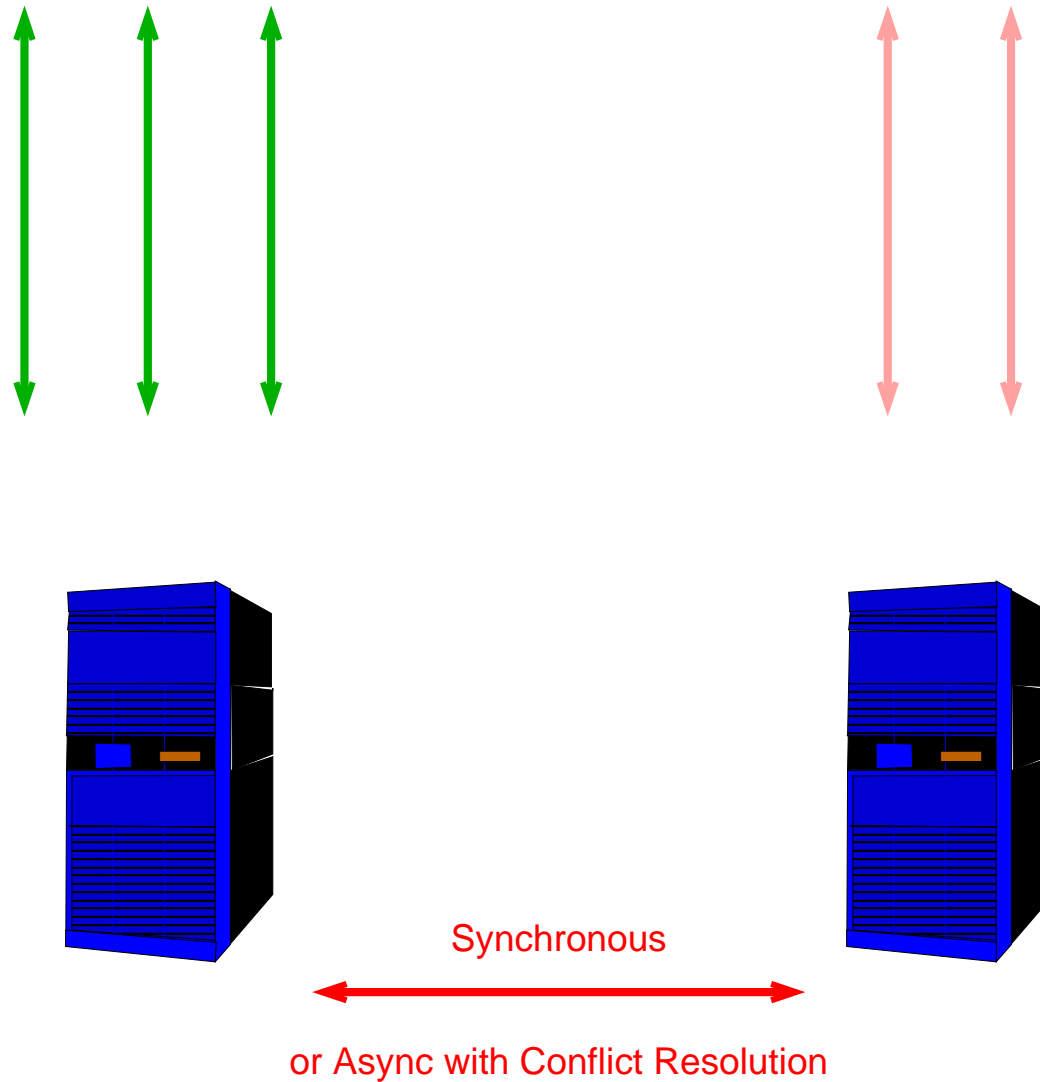


Replication Uses and Methods

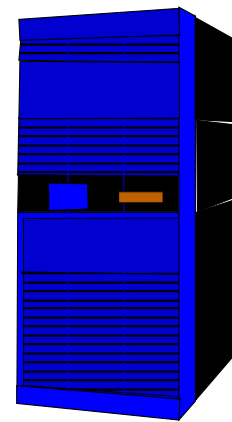
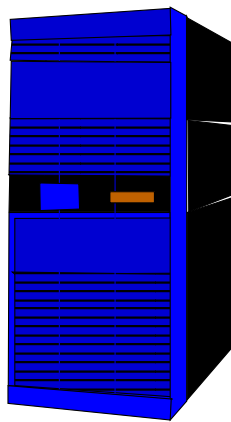
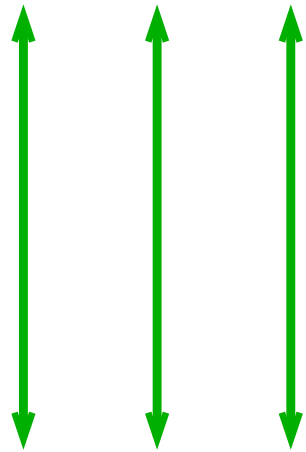
Fail Over



Load Balancing



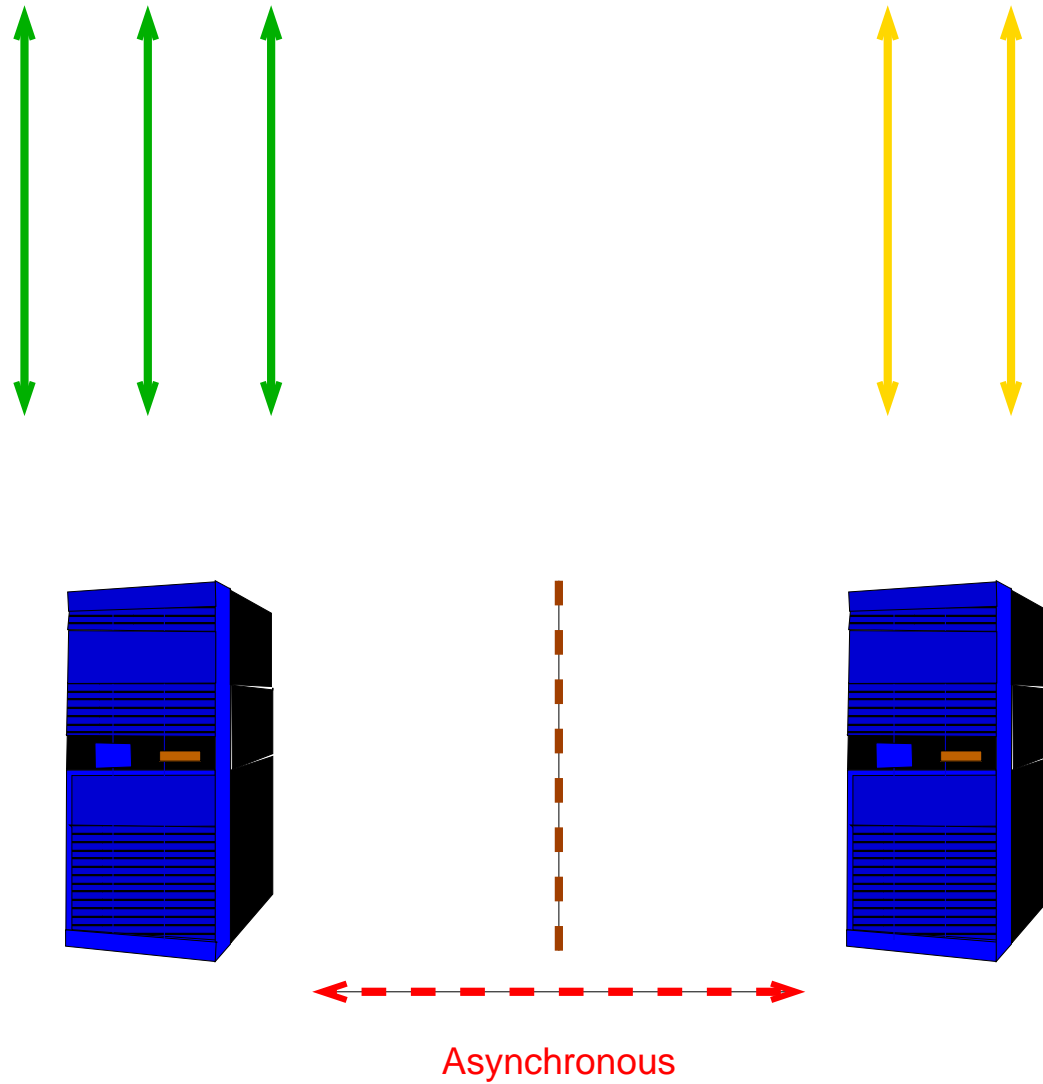
Data Warehousing



Asynchronous

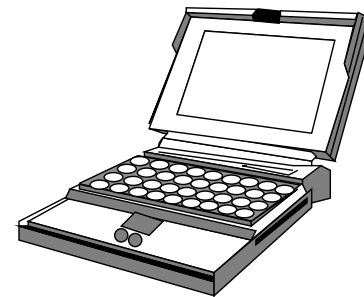
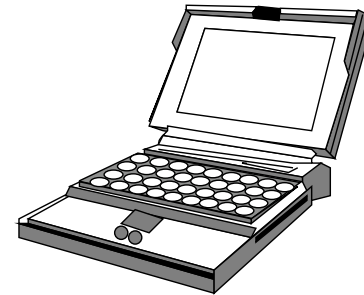
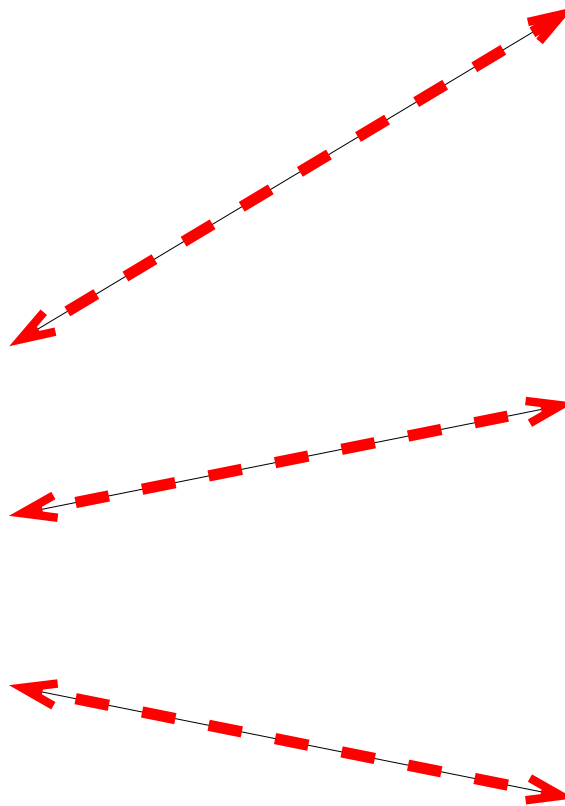
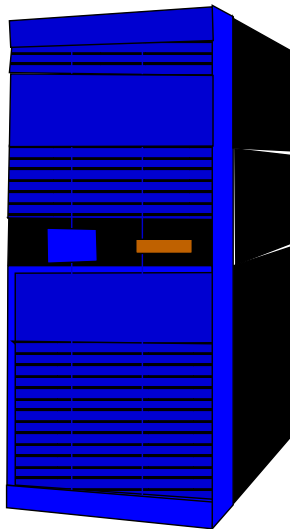


Remote Servers



Mobile Servers

Asynchronous



Similarities

- Data warehousing is read-only load balancing
- Mobile servers are remote servers that connect periodically

Additional Characteristics

- Propagation method
 - SQL
 - Internal query representation
- Replication Event
 - Server process
 - Trigger
 - Log
- Scope
 - Full
 - Partial

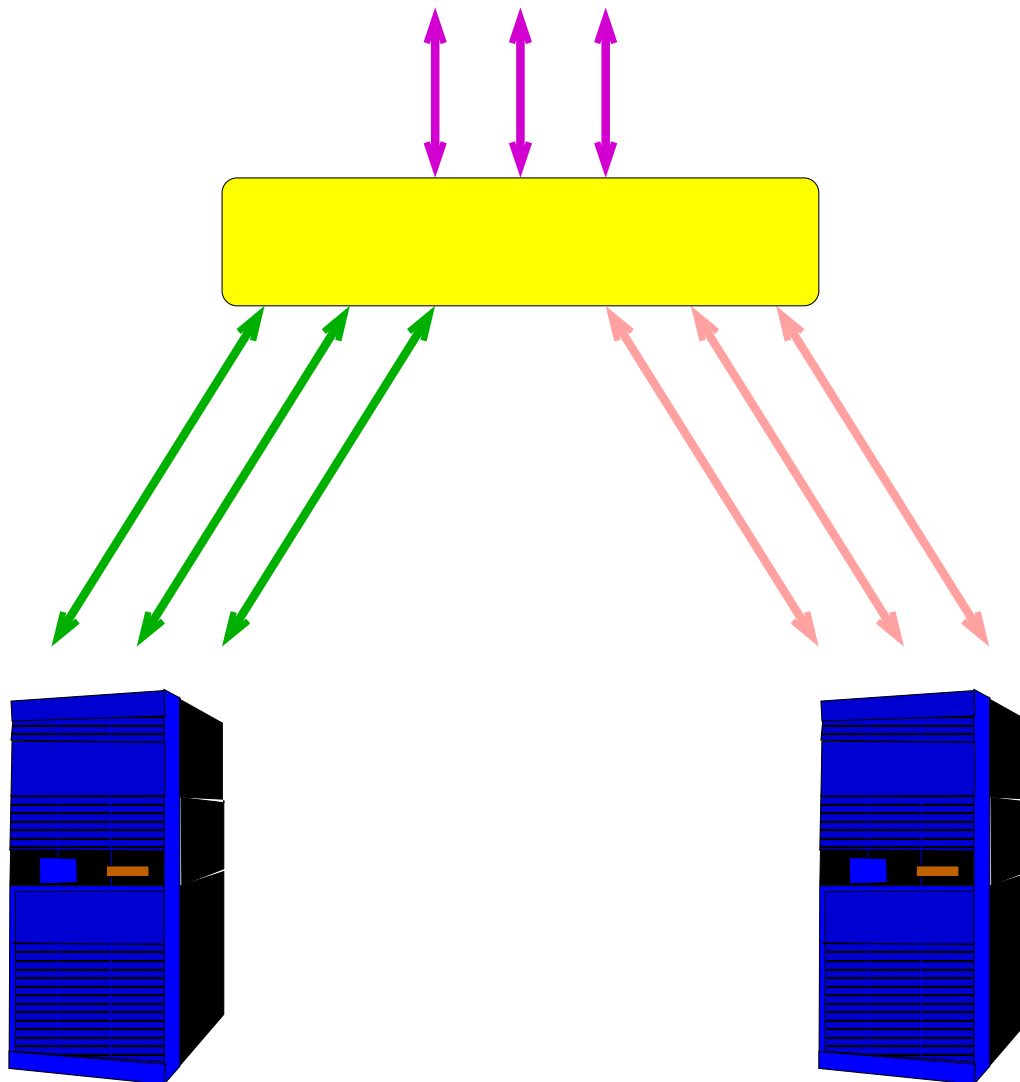
- Integrated
- External

Other Methods

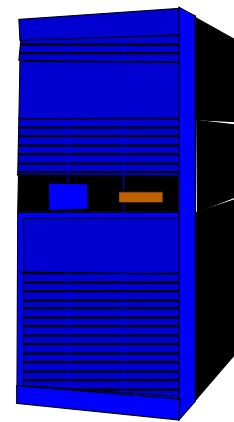
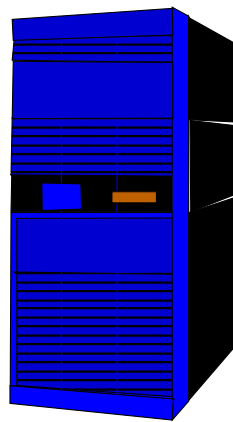
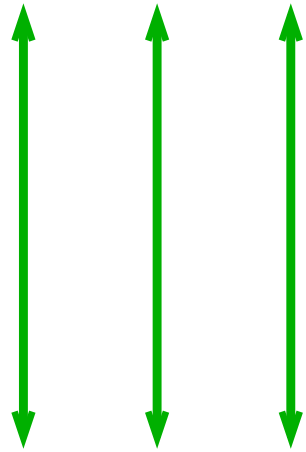
- Multi-ported disks
- Application-level replication
- Data partitioning among servers

Replication Solutions

Usogres



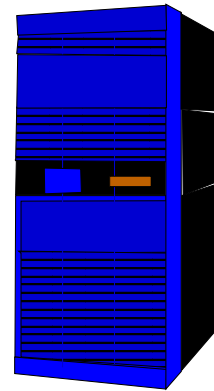
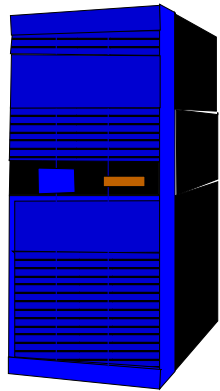
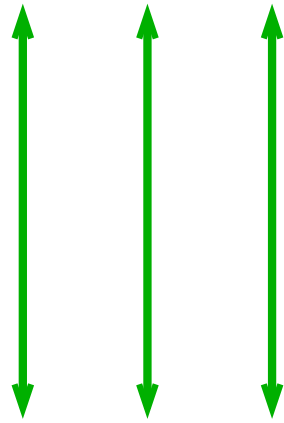
eRServer/Rserv/Dbmirror



Asynchronous



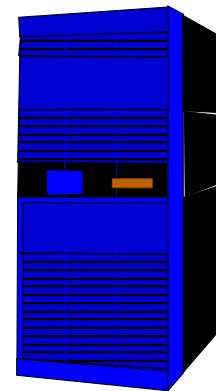
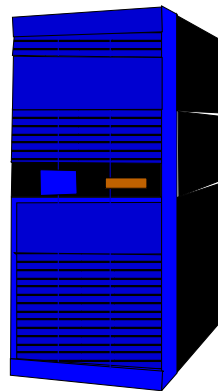
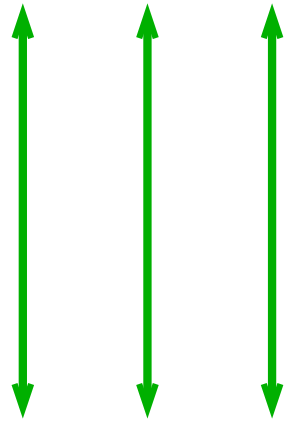
PgReplicator



Asynchronous
with Conflict Resolution

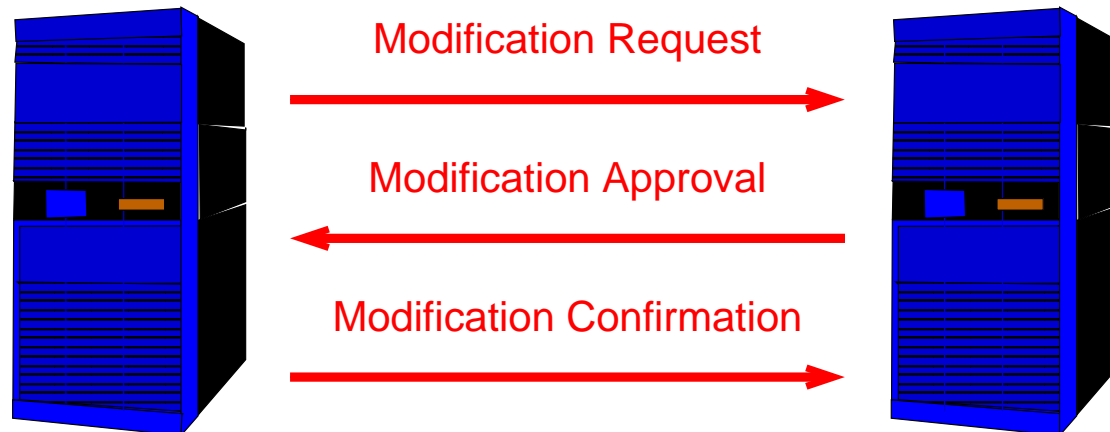
A red dashed double-headed horizontal arrow is positioned below the text, indicating asynchronous replication with conflict resolution between the two nodes.

Postgres-R

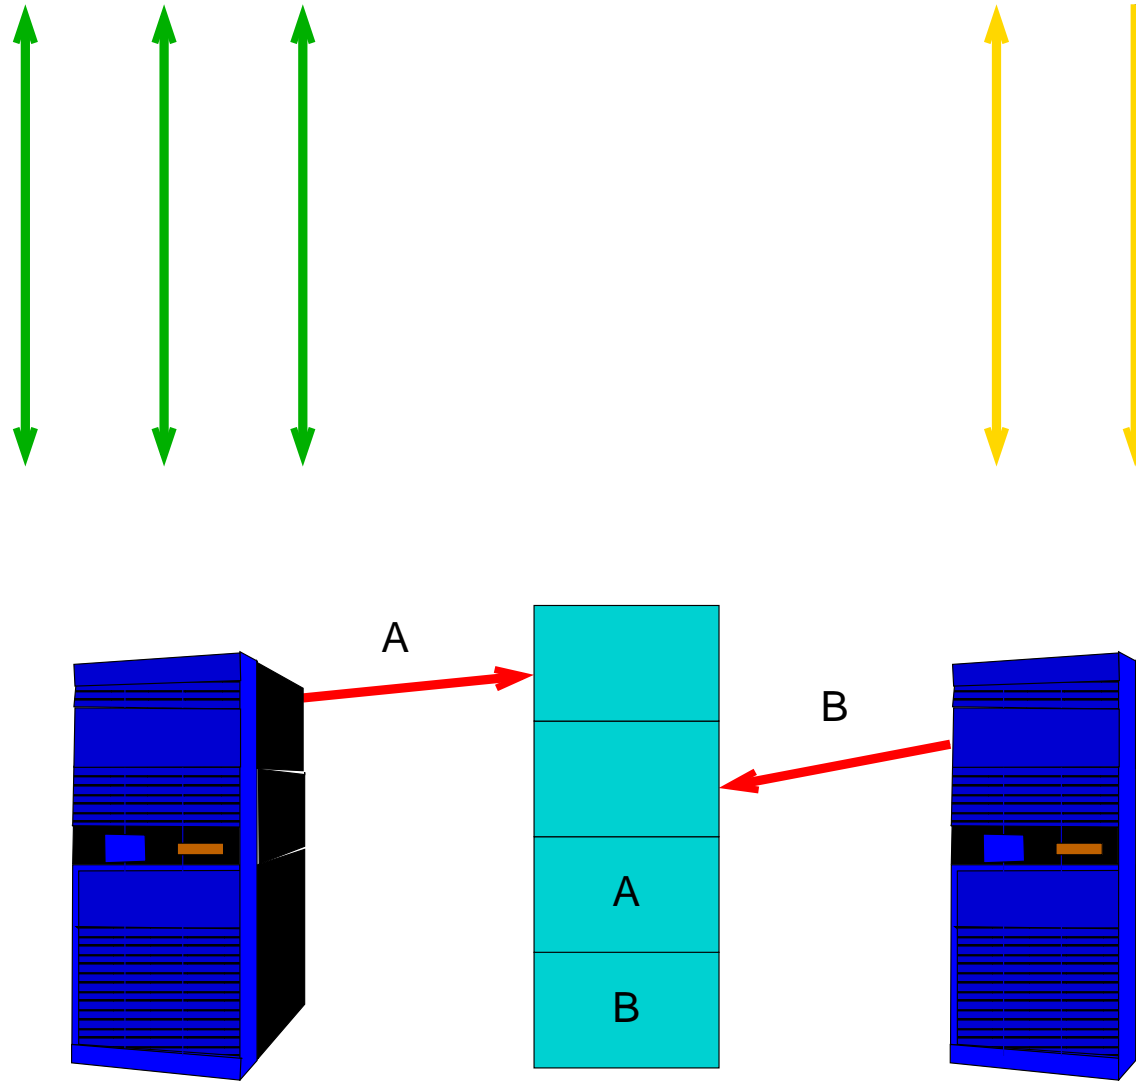


Asynchronous
with No Conflicts

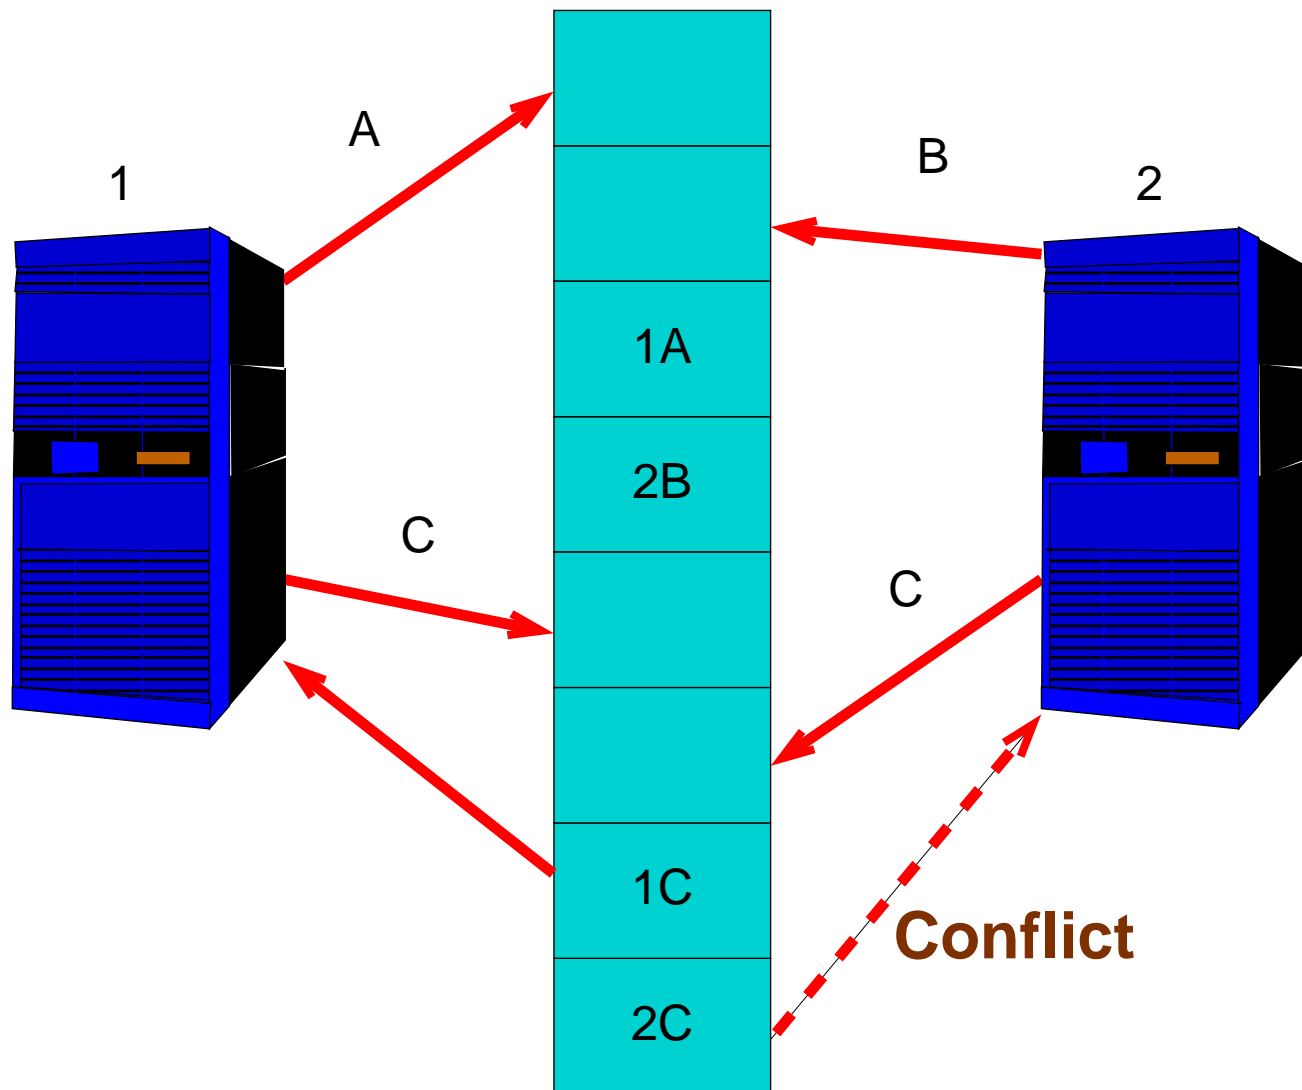
Traditional Synchronization - Two-phase commit



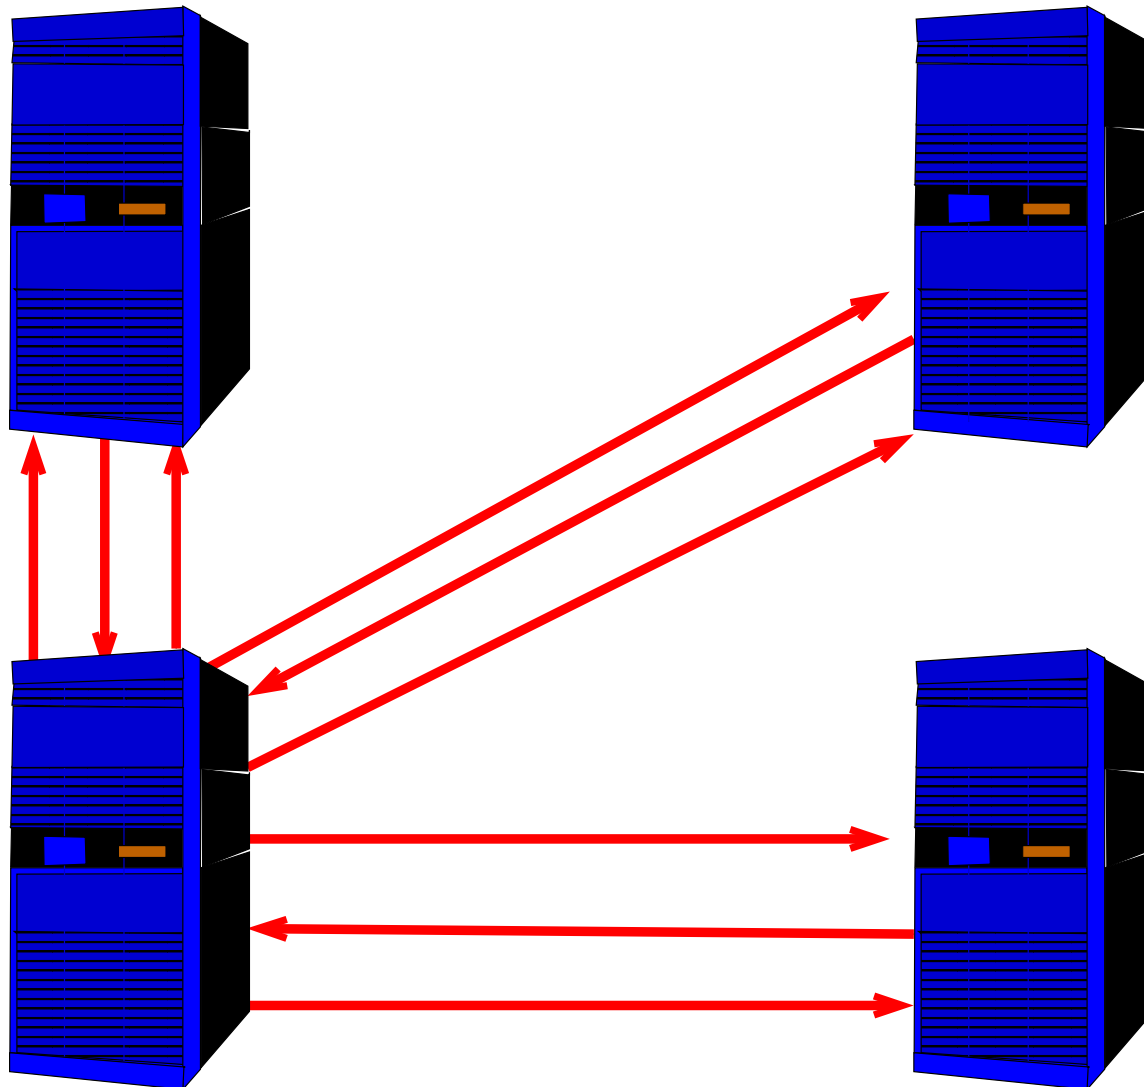
Postgres-R Asynchronization - Group Communication



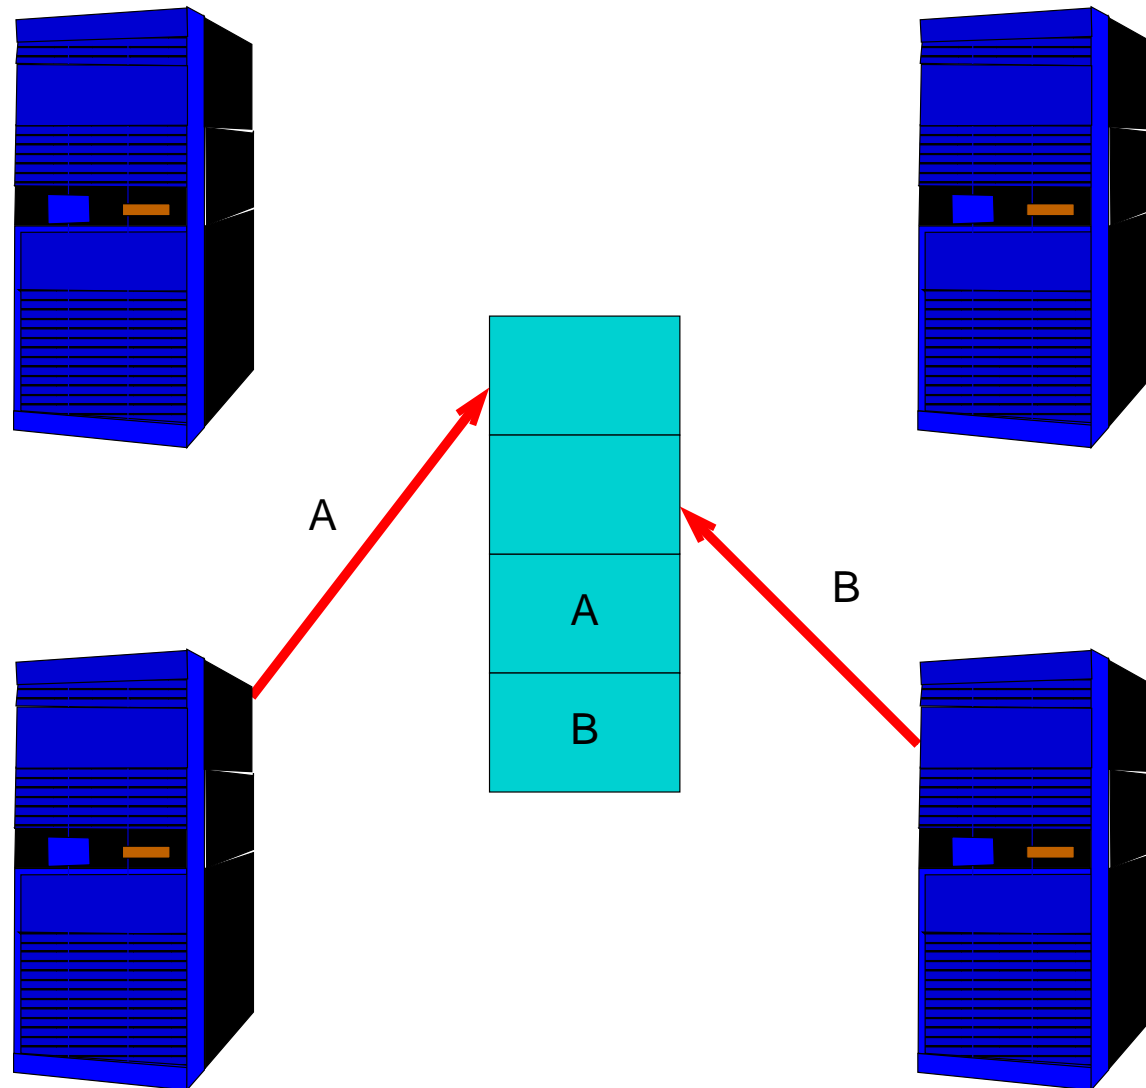
Postgres-R Conflict Detection



Two-phase - Many Servers



Group Communication - Many Servers



Conclusion

